

On the Economic Value and Price-Responsiveness of Ramp-Constrained Storage

Ali Faghih

(Corresponding Author, Email: afaghih@mit.edu)

Massachusetts Institute of Technology (MIT), 77 Mass Ave, Cambridge, MA, USA 02139

Electrical Engineering and Computer Science (EECS) Department, Laboratory for Information and Decision Systems (LIDS)

Mardavij Roozbehani

LIDS, MIT, Cambridge, MA, USA 02139. mardavij@mit.edu

Munther A. Dahleh

EECS, LIDS, MIT, Cambridge, MA, USA 02139. dahleh@mit.edu

The primary concerns of this paper are twofold: to understand the economic value of storage in the presence of ramp constraints and exogenous electricity prices, and to understand the implications of the associated optimal storage management policy on qualitative and quantitative characteristics of storage response to real-time prices. We present an analytic characterization of the optimal policy, along with the associated finite-horizon time-averaged value of storage. We also derive an analytical upperbound on the infinite-horizon time-averaged value of storage. This bound is valid for any achievable realization of prices when the support of the distribution is fixed, and highlights the dependence of the value of storage on ramp constraints and storage capacity. While the value of storage is a non-decreasing function of price volatility, due to the finite ramp rate, the value of storage saturates quickly as the capacity increases, regardless of volatility. To study the implications of the optimal policy, we first present computational experiments that suggest that optimal utilization of storage can, in expectation, induce a considerable amount of price elasticity near the average price, but little or no elasticity far from it. We then present a computational framework for understanding the behavior of storage as a function of price and the amount of stored energy, and for characterization of the buy/sell phase transition region in the price-state plane. Finally, we study the impact of market-based operation of storage on the required reserves, and show that the reserves may need to be expanded to accommodate market-based storage.

Key words: Ramp-Constrained Storage; Dynamic Programming; Inventory Control

1. Introduction

The growing demand for electricity and the urge to reduce greenhouse emissions promote large-scale integration of renewable energy sources, as well as new storage and demand-response technologies to improve energy efficiency in the future grid. However, renewable energy sources are highly uncertain and intermittent. While energy storage technologies can help mitigate the intermittency and narrow the gap between generation from renewable

resources and consumption, they may also add to the uncertainty in the system since the optimal response of storage to prices is a complicated function of both price and the amount of stored energy. Modeling and understanding the behavior of storage in response to real-time market prices is therefore critical for reliable operation of power systems with large amounts of storage.

Energy storage has a clear environmental value; it helps mitigate the intermittency of the renewable resources and thereby maximize their utilization. With sufficient storage capacity, there is no need to curtail generation from renewable energy sources when there is excess of it. Moreover, with proper control policies, storage can help incorporate more renewable resources without compromising reliability. Access to storage also reduces the risk associated with making advance commitments for a renewable generation owner. For instance, if the energy generated by the renewable source falls short of the committed level, the operator can compensate for the shortfall by extracting from the storage. Despite all these potential advantages of storage, if the economic value of storage as an arbitrage mechanism is not attractive, the markets may not invest sufficiently in storage. Hence, unless proper incentives and pricing policies are in place, the environmental and reliability values of storage might not materialize due to underinvestment. Therefore, there is a need for development of econometric models and characterization of the associated optimal policies that can be used for assessing the economic value of storage. This paper seeks to provide such characterization by presenting a model for optimal utilization of ramp-constrained storage in response to stochastically varying electricity prices.

Availability of econometric models of storage and characterization of the effects of storage on the overall price elasticity of demand is also important for system operators who need to maintain stability, and guarantee reliability of the system. It was shown in Roozbehani, Dahleh, and Mitter (2011) that in power grids with information asymmetry between consumers, producers, and system operators, robustness of the system to disturbances is greatly affected by consumers' real-time valuation of electricity, and their response to real-time prices. It was shown that under real-time pricing, high price volatility can be associated with uncertainty in demand response and high price elasticity of demand (PED).

The existing literature covering various dimensions of storage management is extensive. The two major streams of literature on storage that are closely related to this work are the commodity trading/warehousing literature and the electricity storage literature. The warehouse problem is a classical problem in the trading and commercial management of commodities, and has been studied in the literature extensively. As early as 1948, Cahn (1948) introduced the problem of optimizing purchase, i.e. injecting into storage, and sale,

i.e., withdrawing from storage, for the case of a warehouse with fixed size and an initial stock of a certain commodity. Bellman (1956) formulated the warehouse problem in a dynamic programming framework, and Dreyfus (1957) showed that for deterministic prices, the optimal policy at each stage is to either fill up the storage, or empty it, or do nothing, depending on the stage price. Charnes, Drèze, and Miller (1966) solved the warehouse problem for the case of stochastic prices and showed that the optimal policy is the same in the stochastic case (either fill up the storage, or empty it, or do nothing). However, in the warehouse problem studied by these works, there is no limit on the amount that can be injected into, or withdrawn from storage at each stage.

Rempala (1994), and more recently, Secomandi (2010), extended these results for the discrete-time case by imposing a limit on the amount that can be injected into or withdrawn from storage. Kaminski, Feng, and Pang (2008) solved the same extension of the warehouse problem in continuous time. Some other works have also considered this extension of the warehouse problem, including Wu, Wang, and Qin (2011), who propose a heuristic for optimization of seasonal energy storage operations in the presence of ramp constraints, as well as Devalkar, Anupindi, and Sinha (2011) who seek to optimize procurement, processing, and trading of commodities in a multi-period setting. Although the above-mentioned references and our work share similarities in the structure of the optimal policy and the associated value function, the differences in the assumptions on the stochastic price process make the analytical results of these papers different from one another. Unlike all the previous works in this area, we assume that the price at each time period is independent of previous prices. Under this assumption we derive explicit formulas (recursive and/or closed-form) for the thresholds of the optimal policy and for the economic value of storage. We justify our assumption on prices by testing the performance of our optimal policy against price data from real-time markets of the PJM Interconnection (2011) and ISO New England (2011). Another justification for this assumption is that in practice, empirical estimation of conditional distributions (needed for a Markovian price model) requires significant amounts of data. Back of the envelop calculations show that collecting this much data would require going too far back in the price history. However, due to non-stationarity, doing so would make the data irrelevant. Although one can resort to calibrated models for estimation of correlation in data, or try to learn the thresholds directly, we do not pursue these directions in this paper. In addition, we derive an upperbound on the optimal average value per stage of storage in the infinite-horizon case, explicitly highlighting its dependence on storage capacity and the ramp rate. Also different from previous works, we present a computational framework for estimating the average PED

obtained from averaging out the stage-dependence and internal state of storage. We extend our study of the PED induced by storage by addressing the PED as a function of the internal state of storage. This is particularly important in the context of electricity markets because the aggregate price elasticity of demand can affect price volatility and sensitivity to disturbances (Roozbehani, Dahleh, and Mitter 2011). We also highlight another important aspect of state dependence of storage response by showing that under certain assumptions, if the system operator does not know the exact state of storage, the reserves may need to be expanded to accommodate market-based operation of storage.

The literature focusing particularly on electricity storage is also extensive. Bannister and Kaye (1991) studied optimal management of a single storage connected to a general linear memoryless system in the presence of ramp constraints. However, in their model, the objective function is deterministic and the cost is known a priori. Also, Lee and Chen (1995) considered industrial consumers with time-of-use rates and used dynamic programming to determine optimal contracts and optimal sizes of battery storage systems for such consumers. Their work, like ours, pays special attention to the economic value of storage; however, they use a deterministic approach and relax ramp constraints. Several other works have studied the impacts of energy storage on the economics of integration of renewable sources. A renewable generation owner can connect a storage device to the renewable source to optimize the overall profit over time by deciding how much energy to commit to sell at each stage in the time horizon. Some more recent works such as Brown, Lopes, and Matos (2008), Gonzalez et al. (2008), and Korpaas, Holen, and Hildrum (2003) approach this problem by deterministically solving this optimization problem for particular finite sample paths and then averaging the results of these paths. Their approach, as mentioned in Harsha and Dahleh (2011) and Kim and Powell (2011), does not give an optimal policy that can be used in practice because their policy depends on the sample path.

On the other hand, some recent works such as Bitar et al. (2010), Harsha and Dahleh (2011), and Kim and Powell (2011) use dynamic programming to address the problem of managing the revenue of a renewable generator using storage. In particular, Harsha and Dahleh (2011) use a stochastic dynamic programming approach to study the optimal storage investment problem through characterization of optimal sizing of energy storage for efficient integration of renewable resources. In contrast to Harsha and Dahleh (2011), we explicitly include ramp constraints in our model, and highlight the effects of ramp constraints on the value of storage. Another contribution in this line of research is due to Kim and Powell (2011), who study the problem of making advance commitments for a wind generator in the presence of storage capacity and conversion losses. Kim and Powell

(2011) use dynamic programming to obtain the optimal commitment policy when the storage device can have conversion losses but under the assumption of uniformly distributed generation from the wind farm. This assumption allows them to derive the stationary distribution of the storage level and use it to characterize the economic value of storage. In contrast, our results characterize the value of storage purely as an arbitrage mechanism that interacts only with the main grid with a guaranteed supply, without any particular assumption on price distribution other than the assumption of independent prices.

The storage problem in this paper also has similarities with the inventory management literature in terms of the underlying dynamic programming problem and the corresponding optimal policy (for instance, see Federgruen and Zipkin 1986, Kapuscinski and Tayur 1998, and Goel and Gutierrez 2009). However, the main difference comes from the fact that the inventory literature focus on optimally managing inventory in the presence of demand, while the model in this paper assumes that the main grid buys all the energy that we decide to supply, and instead, focuses on trading inventory under ramp constraints and maximizing profit by taking advantage of the fluctuations in the spot prices. Our setup also has similarities with the literature on reservoir hydroelectric system management (see, for instance, Drouin et al. 1996 and Lamond, Monroe, and Sobel 1995) especially those with price uncertainty. However, since the focus of our paper is on the economic value and price-responsiveness of electricity storage as an arbitrage mechanism under ramp constraints, the formulation of our model and our results are different from this line of research.

The model set forth in this paper and some preliminary results were reported in Faghih, Roozbehani, and Dahleh (2011). This paper provides a comprehensive exposition which adds several new ideas, core analytical results, and systematic computational experiments. The contributions of this paper are summarized as follows: First, we propose a dynamic model for optimal utilization of storage in the presence of ramp constraints under the assumption of independent and exogenous prices. This model assumes limited storage capacity and allows sell-back of energy to the grid. Using the principles of stochastic dynamic programming, we analytically characterize the optimal policy and the corresponding value function for the finite-horizon case. In particular, we provide recursive equations for computation of the exact value of storage for the finite-horizon storage problem. To verify the validity of our assumptions, we apply our finite-horizon optimal policy to real-time price data taken from the PJM Interconnection and the Independent System Operator (ISO) of New England. We define the *competitive ratio* (CR) as the ratio of the value obtained from our optimal policy to the absolute maximum value that would have been obtained deterministically, had we known the entire price process a priori. We then show

that the value obtained from our policy under the assumption of independent prices yields a relatively high CR when applied to real-world price data, reaching a CR of about 90% in some cases. This suggests that the sensitivity of the optimal policy and value of storage to the assumption of independent prices is low. We then obtain a closed-form upperbound on the infinite-horizon optimal average value per stage of storage over all possible realizations of prices within a bounded support. To the best of our knowledge, this paper is the first to derive an analytical upperbound on the long-term average value of ramp-constrained storage. This result highlights how the capacity limit and the ramp constraint bound the value of storage. Next, we show that while the economic value of storage is a non-decreasing function of price volatility, the value of storage saturates quickly due to finite ramping rates as the capacity increases, regardless of price volatility. Our results on the economic value of storage can be used to evaluate certain decisions about investment in storage.

Next, in order to study the average price elasticity of a storage system in an electricity market, we study the average PED in a simulated electricity market over a one-day time-horizon, where the term “average PED” reflects the fact that the dependence of storage response on the stage and the internal state of storage has been averaged out. We show that optimal utilization of storage may, in expectation, induce a considerable amount of price elasticity near the average price, but little or no elasticity elsewhere. While the demand for electricity has often been considered to be highly inelastic, the existing literature on price elasticity are mostly based on empirical evidence and qualitative reasoning, see, for instance, Kirschen et al. (2000), Kirschen (2003), Yusta and Dominguez (2002), and Faruqui and George (2002). In this paper, we study price elasticity in a quantitative framework. To the best of our knowledge, this paper is the first to characterize the PED induced by storage through an input-output model of response to prices based on optimal control policies in the presence of ramp constraints. Finally, to examine how the storage response would have been had we not averaged out the state, we characterize the response of storage to exogenous prices and its dependence on the internal state of storage, and highlight the interplay between state-dependence and price-dependence of the response in a computational framework. To eliminate time-dependence, we study these relations in an infinite-horizon setting. We use policy iteration to characterize price responsiveness of a storage system over an infinite time-horizon as a function of the storage state, and characterize the buy/sell phase transition region in the price-state plane. To highlight an interesting implication of the price-state interplay, we study the impact of market-based operation of storage on the required reserves, and show that if the ISO does not have

perfect information about the exact value of the storage state, the reserves may need to be expanded to accommodate market-based operation of storage.

The remainder of this paper is organized as follows: In Section 2, we introduce the dynamic model of utilization of storage. In Section 3, we present the optimal policies for the storage management problem and the corresponding value function, and give an evaluation of the performance of the optimal policy. Next, in Section 4 we derive an analytical upperbound on the optimal average value per stage of storage, and report our computational findings on the economic value of storage. We then discuss the implications of the optimal policy in Section 5, by studying the average PED of a storage system in Subsection 5.1, the price responsiveness of a storage system as a function of the storage state in Subsection 5.2, and the impact of market-based operation of storage on the required reserves in Subsection 5.3. We conclude in Section 6.

2. A Dynamic Model of Storage

2.1. Notation

The set of positive real numbers (integers) is denoted by \mathbb{R}_+ (\mathbb{Z}_+), and non-negative real numbers (integers) by $\overline{\mathbb{R}}_+$ ($\overline{\mathbb{Z}}_+$). The probability mass function (PMF) of a random variable Λ is denoted by P_Λ , and the cumulative distribution function (CDF) is denoted by F_Λ . We will simply use P and F when there is no ambiguity.

2.2. The Model

In this section, we develop a dynamic model for optimal management of ramp-constrained storage in the presence of stochastically-varying prices. We start by formulating the storage management problem as a stochastic dynamic programming problem over a finite horizon.

The Decisions: The decision set of the storage owner or consumer at each discrete instant of time $k \in \overline{\mathbb{Z}}_+$ is characterized by a pair

$$(v_k^{\text{in}}, v_k^{\text{out}}) \in [0, \bar{v}^{\text{in}}] \times [0, \bar{v}^{\text{out}}] \quad (1)$$

where, v_k^{in} and v_k^{out} are, respectively, the amount of power that the consumer injects in, or withdraws from the storage. The corresponding upper bounds (\bar{v}^{in} and \bar{v}^{out}) represent the physical ramp constraints on storage. Also, $v_k = v_k^{\text{in}} - v_k^{\text{out}} \in [-\bar{v}^{\text{out}}, \bar{v}^{\text{in}}]$ denotes the net storage response. With a slight abuse of terminology, we may also refer to the storage response v_k as demand, with the understanding that $v_k \leq 0$ implies a negative demand.

The Price: The price at each stage λ_k is sampled from an exogenous stochastic process that is independently distributed across time, with mean $\bar{\lambda}_k$ and standard deviation σ_k , and support over $[\lambda_k^{\min}, \lambda_k^{\max}] \subset [0, \infty)$. It is assumed that at the beginning of each time interval $[k, k+1]$, the random variable λ_k is materialized and revealed to the consumer. Note that the distributions of λ_k can be different at different stages; however, we assume that the price distribution at each stage is known a priori. We also assume that the feed-in and usage tariffs are the same, i.e., λ_k is the price per unit for both purchase (corresponding to $v_k \geq 0$) and sell-back (corresponding to $v_k \leq 0$), and there are no transaction costs.

The States: The storage state is characterized by a variable

$$s_k \in [0, \bar{s}] \quad (2)$$

where s_k is the amount of energy stored, and \bar{s} is the upper bound on storage capacity. The state s_k evolves according to:

$$s_{k+1} = \beta s_k + \eta^{\text{in}} v_k^{\text{in}} - \eta^{\text{out}} v_k^{\text{out}} \quad (3)$$

where $\beta \leq 1$ is the decay factor, $\eta^{\text{in}} \leq 1$ and $\eta^{\text{out}} \geq 1$ are charging and discharging efficiency factors. Note that efficiency factors and ramp rates might in general be complicated functions of the operating point, i.e., the storage level, but in this work we focus our attention on an ideal case. The idealized model of the dynamics of storage can be written as:

$$s_{k+1} = s_k + v_k, \quad v_k \in [-\bar{v}_{\text{out}}, \bar{v}_{\text{in}}] \quad (4)$$

which corresponds to $\beta = 1$, $\eta^{\text{in}} = 1$, and $\eta^{\text{out}} = 1$.

Penalty and Salvage Value: There is a penalty $h_k(s_k)$ associated with storage, where the sequence of functions $h_k : \mathbb{R}_+ \mapsto \mathbb{R}_+$ are assumed to be monotonic. Also, a salvage value of $\hat{t} \in [\lambda_N^{\min}, \lambda_N^{\max}]$ is assigned to each unit of energy left in storage by the end of the time-horizon.

The Optimization-Based Model of Ideal Storage: Since our goal in this paper is to develop tractable models (and derive bounds) that effectively highlight the important structural features of the optimal control law and the associated economic value of storage with an emphasis on the ramp constraint and storage capacity, we will adopt the idealized model of storage. Nevertheless, in terms of methodology, it would be straightforward to include a “price-adjustment” factor in our formulation to account for injection-withdrawal losses in a similar manner done in Secomandi (2010). In addition, the piecewise-linear penalty function that we have embedded into our model can be used as a surrogate cost

to model the dissipation losses associated with keeping energy in storage. Note also that according to Qian et al. (2011), the academic state of the art is around 95% efficiency for a battery pack and 93% for the overall system with converters. The industrial state of the art is around 90% efficiency for batteries (A123 Systems 2012). We also assume that the ramp constraint is symmetric, i.e. $\bar{v}^{\text{in}} = \bar{v}^{\text{out}} = \bar{v}$. The idealized storage management problem can be formulated as a finite-horizon dynamic programming problem as follows:

$$\begin{aligned}
\min \mathbf{E} & \left[\sum_{k=0}^{N-1} \{h_k(s_k) + \lambda_k v_k\} - \hat{t}s_N \right] \\
\text{s.t. } & s_{k+1} = s_k + v_k \\
& s_k \in [0, \bar{s}] \\
& v_k \in [-\bar{v}, \bar{v}] \\
& \lambda_k \text{ exogenous, and independently} \\
& \text{distributed according to a PMF } P_k
\end{aligned} \tag{5}$$

REMARK 1. We first formulate and solve the storage problem for the finite-horizon case. Later in Section 4.1, we consider the infinite-horizon case and obtain an upperbound on the optimal average value per stage of storage.

3. Characterization of the Optimal Policy and an Evaluation of Its Performance

3.1. The Optimal Policy

In this subsection, we characterize the optimal policy and the value function for problem (5) based on principles of stochastic dynamic programming.

Definition 1. Given a sequence of probability mass functions P_k , $k = 0, 1, \dots, N$, let Θ_k and ψ_k be sequences of maps from the set of all subsets of $\overline{\mathbb{R}}_+$ to $\overline{\mathbb{R}}_+$, defined as follows:

$$\Theta_k : I \mapsto \sum_{\theta \in I} \theta P_k(\theta), \quad \forall I \subset \overline{\mathbb{R}}_+ \tag{6}$$

$$\psi_k : I \mapsto \sum_{\theta \in I} P_k(\theta), \quad \forall I \subset \overline{\mathbb{R}}_+. \tag{7}$$

Given $\bar{v} \in \mathbb{R}_+$, and maps Θ_k and ψ_k as defined in (6) and (7), $\Phi_k^{\bar{v}}$ is the map from the set of all subsets of $\overline{\mathbb{R}}_+$ to \mathbb{R} defined according to

$$\Phi_k^{\bar{v}} : I \mapsto \bar{v}(\Theta_k - \rho\psi_k)I, \quad \forall I \subset \overline{\mathbb{R}}_+, \tag{8}$$

where

$$\rho = \inf I.$$

For instance, Φ_k^1 maps an interval (a, b) to $(\Theta_k - a\psi_k)(a, b)$.

THEOREM 1. Consider the finite-horizon storage management problem (5) with $\bar{s} = n\bar{v}$ for some $n \in \mathbb{Z}_+$. Furthermore, assume that the penalty functions $h_k : [0, \infty) \rightarrow [0, \infty)$, $k = 0, \dots, N$ are piecewise linear non-decreasing convex functions of the form:

$$h_k(s) = h_k^i s + c_k^i, \quad s \in [i\bar{v}, (i+1)\bar{v}), \quad i \in \overline{\mathbb{Z}}_+ \quad (9)$$

Then,

(i) the optimal policy is characterized as follows:

if $s_k \in [i\bar{v}, (i+1)\bar{v})$, for some $i \in \overline{\mathbb{Z}}_+$, then

$$v_k^* = \begin{cases} \max(-s_k, -\bar{v}), & t_{k+1}^{\max(0, i-1)} < \lambda_k \text{ for all } i \\ i\bar{v} - s_k & , t_{k+1}^i < \lambda_k \leq t_{k+1}^{i-1} \text{ for } i \geq 1 \\ (i+1)\bar{v} - s_k & , t_{k+1}^{i+1} < \lambda_k \leq t_{k+1}^i \text{ for all } i \\ \bar{v} & , \lambda_k \leq t_{k+1}^{i+1} \text{ for all } i \end{cases} \quad (10)$$

where the thresholds are computed via the following recursive equations:

$$\begin{aligned} t_N^i &= \hat{t}, & i \in \{0, 1, 2, \dots, n-1\} \\ t_N^i &= -h_N^i, & i \geq n \\ \text{for } k < N : & \\ t_k^0 &= t_{k+1}^1 + \Phi_k^1(t_{k+1}^1, \lambda_k^{\max}) - h_k^0 \\ t_k^i &= t_{k+1}^{i-1} - h_k^i + \Phi_k^1(t_{k+1}^{i+1}, t_{k+1}^{i-1}) + (t_{k+1}^{i+1} - t_{k+1}^{i-1})F_k(t_{k+1}^{i-1}), & i \geq 1 \end{aligned} \quad (11)$$

(ii) the value function is a piecewise linear convex function of the form:

$$V_k(s) = -t_k^i s + e_k^i, \quad s \in [i\bar{v}, (i+1)\bar{v}), \quad i \in \overline{\mathbb{Z}}_+ \quad (12)$$

where t_k^i are the thresholds given in (11) and the intercepts e_k^i are computed via the following recursive equations:

$$\begin{aligned} e_N^i &= 0, & i \in \{0, 1, 2, \dots, n-1\} \\ e_N^i &= \bar{s}(t_N^i - \hat{t}), & i \geq n \\ \text{for } k < N : & \\ e_k^0 &= c_k^0 + e_{k+1}^0 + (\bar{v}\lambda_k^{\min} + e_{k+1}^1 - e_{k+1}^0 - \bar{v}t_{k+1}^1)F_k(t_{k+1}^0) \\ e_k^i &= c_k^i + \bar{v}\lambda + f(t_{k+1}^{i-1}, t_{k+1}^i, t_{k+1}^{i+1}, e_{k+1}^{i-1}, e_{k+1}^i, e_{k+1}^{i+1}) + g(t_{k+1}^{i-1}, t_{k+1}^i, t_{k+1}^{i+1}), & i \geq 1 \end{aligned} \quad (13)$$

where the functions f and g are given by

$$f(\cdot) = e_{k+1}^{i-1} - \bar{v}t_{k+1}^{i+1} + (e_{k+1}^i - e_{k+1}^{i-1})F_k(t_{k+1}^{i-1}) + (e_{k+1}^{i+1} - e_{k+1}^i)F_k(t_{k+1}^i)$$

$$g(\cdot) = (i+1)\Phi_k^{\bar{v}}(t_{k+1}^{i+1}, t_{k+1}^i] + i\Phi_k^{\bar{v}}(t_{k+1}^i, t_{k+1}^{i-1}] - \Phi_k^{\bar{v}}(t_{k+1}^{i-1}, \lambda_k^{\max}] - \Phi_k^{\bar{v}}(t_{k+1}^i, \lambda_k^{\max}].$$

Proof. Please see the Appendix. \square

Figure 1 shows how the thresholds vary with time and state for the case of a discretized truncated log-normal distribution with mean $\bar{\lambda}_k = 49$ and $\sigma_k = 9$ for all k , i.e. for independently and identically distributed (i.i.d) prices. For generating this plot, we set $\bar{v} = 1$, $n = 15$, $N = 24$, $\hat{t} = \bar{\lambda}_k$, and assume no storage penalties (i.e. $h_k^i = 0$ for $i < n$ and all k).

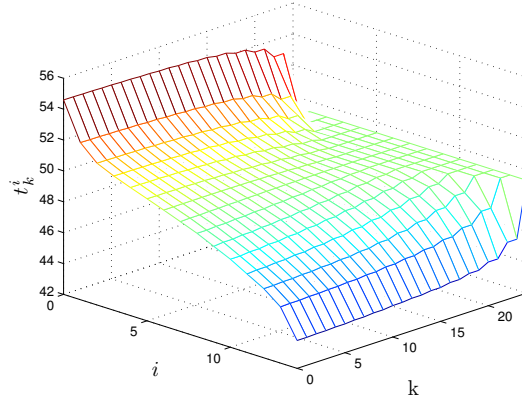


Figure 1 Thresholds as a function of time and state for a discretized truncated log-normal distribution (and i.i.d prices) with mean $\bar{\lambda}_k = 49$ and $\sigma_k = 9$ for all k , $\bar{v} = 1$, $n=15$, $N = 24$, $\hat{t} = \bar{\lambda}_k$ and no storage penalties.

REMARK 2. The form of the optimal policy shows that if we start with an empty storage ($s_0 = 0$), then the storage state s_k will only take integer multiples of \bar{v} since $v_k^* \in \{-\bar{v}, 0, \bar{v}\}$ for all k . When $s_0 \neq 0$, the storage state will fall on the grid of integer multiples of \bar{v} immediately after the first time that $v_k^* = i\bar{v} - s_k$ or $v_k^* = (i+1)\bar{v} - s_k$, and we will have $v_k^* \in \{-\bar{v}, 0, \bar{v}\}$ for the remainder of the time horizon. This conclusion holds for the infinite-horizon case as well and can help simplifying the policy and analysis by focusing on a finite state system.

REMARK 3. The upperbound \bar{s} on the storage capacity is enforced by choosing h_k^i in (9) sufficiently large (i.e. $h_k^i > \lambda_k^{\max}$) for $i \geq n$, so that it would never be optimal to store energy beyond \bar{s} . It can be verified that the thresholds and consequently the optimal policy are invariant with respect to the choice of h_k^i for $i \geq n$ as long as $h_k^i > \lambda_k^{\max}$.

Definition 2. We define the economic value of storage, or simply the value of storage, as the negative of the cost of the optimal value of problem (5) divided by the number of stages (N), and we denote it by \mathcal{V} for the finite-horizon case and by \mathcal{V}_∞ for the infinite-horizon

case. Therefore, $\mathcal{V} = -V_0(s_0)/N$. For instance, if $s_0 = 0$ (i.e. the consumer starts with an empty storage), it then follows from (12) that for the finite-horizon case:

$$\mathcal{V} = -V_0(0) = -e_0^0/N. \quad (14)$$

Thus, \mathcal{V} can be computed using the recursive equations in (13) for the finite-horizon case.

3.2. The finite-horizon value of storage under empirical price distribution from real-time wholesale electricity markets

Herein, to examine our results and verify the validity of our assumptions, we test the optimal policy against actual market price data. We apply the finite-horizon optimal policy (10) to real-time wholesale market data taken from the PJM Interconnection (2011) and the ISO New England (2011) to examine the competitive ratio (CR) of our policy when applied to real-world price data. This would allow us to assess whether the assumption of independent prices is actually reasonable for the storage problem. A relatively high CR would not necessarily suggest that prices are independent; rather, it could mean that the sensitivity of the optimal policy to the correlations, if any, in real-world prices is low. If the CR obtained from the assumption of independent prices is relatively high, we can suggest that the additional information that the current price provides on future prices (as a conditional distribution) has little value.

The setup of the experiment is as follows. We take the actual data for hourly energy prices (for 16 hours of each day, from 8 a.m. to 12 a.m.) for two different months (December 2010 and July 2011) from the PJM Interconnection (2011), and May and November 2011 from ISO New England (2011). The choice of these dates and times was arbitrary. For the purpose of these simulations, we set \hat{t} , the salvage value, equal to $\bar{\lambda}$, the empirical mean. To perform the simulations, we first find the empirical distribution of the data for each case. In these computations, we use the same price distribution for all k ; i.e., we take all the price data for the entire month, and use this data to find the empirical distribution of prices in that month and compute the thresholds for the optimal policy for all hours from the same empirical distribution. In other words, we assume that prices are independently and identically distributed (i.i.d). Then, at the beginning of each stage we reveal the actual price of that stage to the optimal policy and record the decision. This gives us the profit resulting from the optimal policy. Next, for the purpose of comparison, we assume the entire price sequence is perfectly known a priori and compute the value that results from deterministically and omnisciently maximizing the profit against the materialized prices. This deterministic value is the absolute best that an omniscient agent could have done,

Table 1 Competitive Ratio for each set of simulations

Source/Month	Competitive Ratio	# of qualifying days	Range from the mean
PJM / July	0.86	27	$1 \times \sigma$
	0.84	28	$1.5 \times \sigma$
	0.81	All	All
PJM / December	0.88	22	$1 \times \sigma$
	0.87	25	$1.5 \times \sigma$
	0.77	All	All
ISO NE/ May	0.72	25	$1 \times \sigma$
	0.71	All	$1.5 \times \sigma$
	0.71	All	All
ISO NE/ November	0.89	29	$1 \times \sigma$
	0.88	All	$1.5 \times \sigma$
	0.88	All	All

and we will refer to its corresponding deterministic policy as the omniscient policy. We then find the CR by computing the ratio of the value obtained from the optimal policy to the value obtained from the omniscient policy. The CR gives us a measure of how well our optimal policy has performed.

For each month, there is a number of days in which the average price is well above the average price of the entire month. These days are outliers in the sense that the empirical distribution is way off for modeling their price sequence. We go about this issue by performing three sets of simulations. For the first set of simulations, we only select those days in which the average price is within one standard deviation of the average price of all the selected days in the month. In the second set of simulations, we select those days in which the average price is within about 1.5 times the standard deviation of the average price of all the selected days in the month. Then, in the third set of simulations, we take all the days of the month into account, even the outliers. We record the CR for each day, and report the average CR for each month and each case in Table 1. Comparing the results in Table 1 reveals how much these outlier days affect the CR. Note that a ramp constraint of $\bar{v} = 1$ and a storage capacity of $\bar{s} = 10$ is used in all the simulations, and it is assumed that there are no penalties on storage.

Figure 2 shows the plots of the empirical value of storage for each day of the month in the second set of simulations (i.e for those days whose average price is within 1.5 times the standard deviation of the average price of all the selected days in the month).

As it can be seen in Figure 2, we have nearly perfect matching in some days, while in some other days there is discrepancy. This discrepancy is mainly due to two reasons. First, it is because we have multiple spikes in some days with only a few (or no) prices that are below our buying thresholds. So, even though there are ample opportunities for arbitrage,

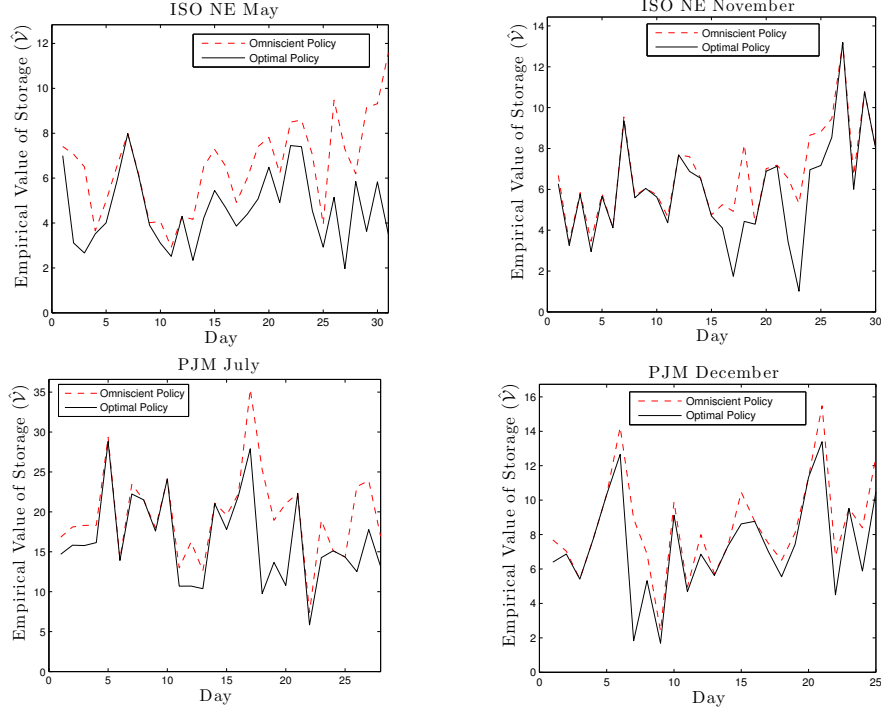


Figure 2 The empirical value of storage for each day of the month in the second set of simulations, obtained from applying the optimal policy (solid line) and the omniscient policy (dashed line)

even the lowest of these spiky prices is not within the normal range. In Table 1, this effect can be observed for the month of December in PJM, in which the removal of those days with high average prices improved the CR from 0.77 to 0.88. The second reason is that in some days all the prices are almost in the same range compared to our thresholds (i.e. they are either mostly above our buying thresholds or mostly below our selling thresholds). In other words, the low CR is just an outcome of an undesirable sequence of prices (sample path). So, even though for the deterministic case with the price sequence known a priori it is possible to take advantage of these small price differentials, our thresholds are unable to capture these opportunities. This effect can be observed in the month of May in Table 1, in which the removal of those days with high average prices did not really improve the CR; this is because the average price in the days with a relatively flat price sequence is not necessarily considerably higher than the month's average.

So far, we have been computing the thresholds for the optimal policy using the empirical price distributions from the prices in that same month. This can be taken as a proxy for the sensitivity of the value of storage to correlations in the actual prices, given that our optimal policy assumes independent prices. Although there is no benchmark to compare with, the CR obtained from applying the optimal policy seems reasonably high, and we suggest that the sensitivity of the value of storage to the assumption of independent prices

is low. Also, although we do not do as well on those days with higher than normal, or very flat price profiles, it does not appear that a Markovian or Martingale assumption on prices could do any better, because these prices are really outliers and do not seem to follow a structured stochastic pattern that can be learned from past data. However, this needs to be substantiated by further studies and systemic experiments.

In the next set of experiments, instead of computing the empirical distribution from the data of that same month, we compute the empirical distributions from the price data of the past 30 days, the past 20 days, and the past 10 days, respectively. The results are shown in Table 2.

Table 2 Competitive ratios using the empirical distribution from the past 30, 20, and 10 days

Source / Month	Past 30 days	Past 20 days	Past 10 days	# of selected days
PJM / July	0.68	0.70	0.75	All
PJM / December	0.55	0.65	0.67	All
ISO NE/ May	0.69	0.70	0.68	All
ISO NE/ November	0.91	0.89	0.89	All

An interesting observation is that for both months of May and November in ISO NE, using the empirical distribution from the past 20 days gives almost the same CR as using the empirical distribution from May and November themselves (as reported in Table 1). However, comparing the competitive ratios shown in Table 2 for both months in PJM with the results shown in Table 1, we observe that for PJM, using the empirical distribution from historical data does not do as well as using the empirical distribution from that month itself.

4. Characterization of the Economic Value of Storage

4.1. Analytical upperbound on the value of storage

In this subsection, we will derive a bound on the long-term economic value of ramp-constrained storage. Herein, for the purpose of obtaining the bound, we further assume that $h_k^i = 0$, for all $i \leq n - 1$ and $k \leq N$. This assumption is consistent with our objective of finding an upperbound on the value of storage.

Definition 3. Given a control policy $\pi_k : [0, \infty)^2 \mapsto [-\bar{v}, \bar{v}]$, and starting from an arbitrary initial state s , the infinite-horizon average cost per stage associated with problem (5) is defined as

$$\gamma_\pi \stackrel{\text{def}}{=} \lim_{N \rightarrow \infty} \frac{1}{N} \mathbf{E} \left[\sum_{k=0}^{N-1} \lambda_k v_k \mid s_0 = s \right], \quad (15)$$

where $v_k = \pi_k(x_k, \lambda_k)$. We will refer to the problem of optimization of γ_π over all feasible stationary policies as the infinite-horizon storage management problem. We will denote the associated optimal cost by γ^* and refer to $\mathcal{V}_\infty \stackrel{\text{def}}{=} -\gamma^*$ as the long-term expected economic value of storage.

REMARK 4. It is standard to show that the optimal average cost is independent of the initial state s_0 . Moreover, if the relative value iteration for the infinite-horizon storage problem converges to some differential cost function $H^*(s)$, then it is necessary for $H^*(s)$ and the optimal average cost per stage γ^* to satisfy the Bellman equation (see, for instance, Bertsekas 2000):

$$H^*(s) = \mathbf{E} \left[\min_{v \in [\max(-s, -\bar{v}), \min(\bar{v}, \bar{s}-s)]} \lambda v + H^*(s+v) \right] - \gamma^*. \quad (16)$$

THEOREM 2. *Consider the infinite-horizon storage management problem. Suppose that the support of the price distribution function at all stages lies within an interval $[\lambda_{\min}, \lambda_{\max}] \subseteq [0, \infty)$. All else held constant, the maximum over all possible distributions, of the long-term economic value of storage is given by*

$$\mathcal{V}_\infty = -\gamma^* = \frac{\bar{v}(\lambda_{\max} - \lambda_{\min})}{2} \frac{n}{n+1} = \frac{(\lambda_{\max} - \lambda_{\min})}{2} \frac{\bar{s} \bar{v}}{\bar{s} + \bar{v}}, \quad (17)$$

and is attained when the prices are sampled from a two-point symmetric distribution with nonzero probability masses placed at the endpoints of the fixed support:

$$P_\Lambda(\lambda) = \begin{cases} 1/2 & \text{if } \lambda = \lambda_{\min} \\ 1/2 & \text{if } \lambda = \lambda_{\max} \\ 0 & \text{otherwise} \end{cases} \quad (18)$$

Furthermore, the corresponding differential cost function satisfying the Bellman equation is given by the following piecewise linear convex function:

$$H^*(s) = -\frac{(i+1)\lambda_{\min} + (n-i)\lambda_{\max}}{n+1} s - \frac{i(i+1)(\lambda_{\max} - \lambda_{\min})\bar{v}}{2(n+1)}, \quad (19)$$

where $s \in [i\bar{v}, (i+1)\bar{v})$ and $i \in \{0, \dots, n-1\}$, and for the special case of $s = \bar{s}$, we take $i = n-1$.

Finally, for any two-point distribution with PMF

$$P_\Lambda(\lambda) = \begin{cases} a & \text{if } \lambda = \lambda_{\max} \\ 1-a & \text{if } \lambda = \lambda_{\min} \\ 0 & \text{otherwise} \end{cases} \quad (20)$$

we have

$$\mathcal{V}_\infty = -\gamma^* = \bar{v}(\lambda_{\max} - \lambda_{\min}) \frac{b(1+b+\dots+b^{n-1})}{(b+1)(1+b+\dots+b^n)} \quad (21)$$

where $b = (1-a)/a$.

Proof. Please see the Appendix. \square

If in addition to the support of the price distribution we also fix the mean of the distribution, we can obtain a tighter bound as stated in Corollary 1:

COROLLARY 1. *Suppose that in addition to fixing the support of the price distribution function, we also fix the mean of the price distribution to $\mu \in (\lambda_{\min}, \lambda_{\max})$. All else held constant, the maximum over all possible distributions, of the long-term expected value of storage is attained when the prices are sampled from a two-point distribution with the following PMF:*

$$P_{\Lambda}(\lambda) = \begin{cases} \frac{\mu - \lambda_{\min}}{\lambda_{\max} - \lambda_{\min}} & \text{if } \lambda = \lambda_{\max} \\ \frac{\lambda_{\max} - \mu}{\lambda_{\max} - \lambda_{\min}} & \text{if } \lambda = \lambda_{\min} \\ 0 & \text{otherwise} \end{cases}$$

The corresponding long-term value of storage is obtained by plugging $b = (\lambda_{\max} - \mu)/(\mu - \lambda_{\min})$ into (21).

Proof. Please see the Appendix. \square

REMARK 5. The $n/(n+1)$ scaling in the optimal average cost per stage implies that 90% of the maximum possible value of storage is achieved when the storage capacity is only 9 times the ramp constraint. Note that this was obtained for an extreme distribution. As we will see in the remainder of this section, for less extreme distributions with smaller variance over the support, the value of storage saturates even more quickly. This includes empirical distributions obtained from electricity market data. Furthermore, aging, dissipation, and non-ideal charging and discharging factors further reduce the value of storage.

4.2. Computational Experiments for Characterization of the Economic Value of Storage

In this subsection, we employ numerical computations to characterize the economic value of the proposed model of storage over a finite time-horizon and highlight the effects of ramp constraints and price volatility on the value of storage.

Herein, we consider the following classes of distributions:

- Discretized truncated log-normal distribution, with fixed mean $\bar{\lambda} = 50$,
- Discretized uniform distribution, with fixed mean $\bar{\lambda} = 50$.

The reason for choosing the log-normal distribution is that the empirical distributions from ISO New England and PJM qualitatively resemble a log-normal distribution, at least for the cases that we tested. The choice of the mean ($\bar{\lambda} = 50$) is also a realistic choice for average hourly energy prices in markets such as PJM and ISO NE. For the purpose of

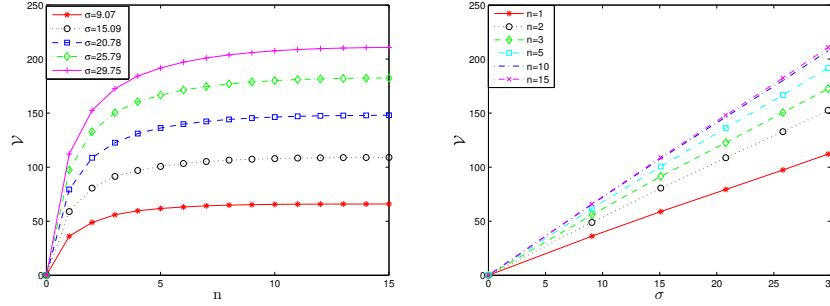


Figure 3 \mathcal{V} vs. n (left) and σ (right) for a few samples, using a discretized truncated log-normal distribution.

these computations, we use the same price distribution for all k (i.e. we assume that prices are independently and identically distributed). Throughout this section we assume that $s_0 = 0$, which means that the consumer starts with an empty storage. For each of these distributions, we fix all quantities in our model other than σ , the standard deviation of price distribution, and n , the ratio of storage capacity \bar{s} to physical ramp constraint of storage \bar{v} . We vary $n = \bar{s}/\bar{v}$ by fixing \bar{v} and changing \bar{s} . Using the fixed quantities $N = 24$, $\bar{v} = 10$, and $\bar{\lambda} = 50$, we examine how \mathcal{V} varies as a function of σ and n for the case of no storage penalties. For the purpose of these simulations, we set \hat{t} equal to the mean of the price distribution. Herein, we set $h_k^i = 0$, for all $i \leq n - 1$ and $k \leq N$, so that there is no penalty on storing energy up to capacity. Then, for a fixed time horizon, we examine how \mathcal{V} varies with σ and n , for each of the following price distributions:

Discretized truncated log-normal distribution: Figure 3 illustrates how \mathcal{V} changes with σ and n , for the discretized truncated log-normal distribution. The plots show that the value of storage increases linearly with σ . As one would expect, the value of storage also increases as the storage capacity increases. However, it is interesting to note that for a fixed standard deviation, the value of storage saturates fairly quickly as a function of n . Hence, for a given time horizon, a fixed ramp constraint, and a fixed σ , there exists a certain range for capacity beyond which the value of storage will no longer change noticeably. Also, the optimal storage capacity increases with price volatility.

Discretized uniform distribution: As can be seen in Figure 4, saturation of the value of storage occurs at about the same value as in the log-normal case. However, for certain extreme distributions, such as an asymmetric 2-point distribution, saturation can occur more quickly. Note also that the value of storage is a linear function of the standard deviation, just like the log-normal case.

One interesting observation in these results is that in the presence of ramp constraints, several distributed storage systems would be more profitable than one large storage system of equal ramp constraint and aggregate capacity, due to the quick saturation of \mathcal{V} as n

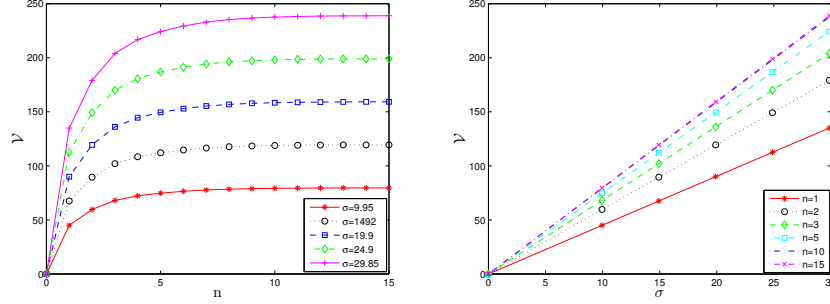


Figure 4 \mathcal{V} vs. n (left) and σ (right) for a few samples, using a discretized uniform distribution.

increases. Although this observation is based on the assumption that the ramp constraint and capacity are independent, this assumption might actually be valid for the case of distribution grids in which the ramp constraint is imposed by the power lines. Another interesting observation is that although the shape of the plots look quite similar for both distributions, the value of the uniform distribution is higher than that of the log-normal. This makes intuitive sense because with the uniform distribution, we have as many opportunities for buying at a low price as we have for selling at a high price, while for the log-normal case, most of the probability mass is centered around the mode, which creates fewer opportunities for arbitrage.

Using wholesale market data from ISO New England: For the purpose of comparison, we repeat the experiment using the data for hourly energy prices from three different months (April, June, and October of 2011) obtained from ISO New England (2011). For each month, the hourly data (from all 24 hours of all days of the month) have been used to find the empirical distribution of prices for that particular month. For the purpose of these computations, we compute the empirical distribution over all stages, and we use this empirical distribution for all the stages (i.e. we assume that prices are i.i.d). Figure 5 shows how \mathcal{V} varies with n in each month without storage penalties:

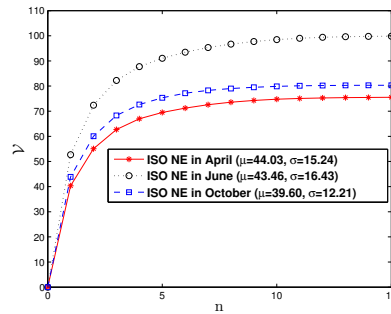


Figure 5 \mathcal{V} vs. n for 3 different months, without storage penalties

Note in Figure 5 that similar to the results of the distributions used in Section 4.2, the value of storage nearly saturates when n is about 6. An interesting observation is that

although both the mean and the standard deviation of the empirical price distribution in October are lower than those of the empirical price distribution in April, the value of storage is higher in October. This observation reiterates the importance of the price distribution for the value of storage, even when its mean and standard deviation are somewhat lower, though many other factors such as non-stationarity and correlation in the data could skew the results.

5. Implications of the Optimal Policy

Herein, we will focus on understanding the price-dependence and state-dependence of the optimal policy, and the impact of the state-price interplay on the storage response. We first average out state-dependence of storage response and focus our attention on “average” price-responsiveness only, which allows us to address the expected price elasticity of demand induced by storage in the first subsection. Then, in the next two subsections, we focus on the interplay between state-dependence and price-dependence of storage response, and study potential impacts of the state-and-price-dependent response from market-based storage on the required reserves.

5.1. Average Price Elasticity of Demand Induced by Storage

In this subsection, in order to study the implications of the optimal policy on price elasticity of demand (PED) in an electricity market, we present a computational framework for studying the average PED in a simulated electricity market. In our dynamic model, the storage response depends on price, stage, state, time-horizon, storage capacity, and ramp constraint. The term “average PED” reflects the fact that the dependence of storage response on the stage and the internal state of storage have been averaged out. We assume that there is a fixed time horizon N , and eliminate state-dependence by taking expectations. In particular, we define:

$$v(k, \lambda) = \mathbf{E}_{s_0, s_k} [v_k^* | \lambda_k = \lambda].$$

In order to eliminate stage-dependence, we think of the storage response-measuring observer as sampling a random time τ uniformly over $\{0, \dots, N\}$. By averaging over this randomness, we maintain dependence on price alone:

$$v_{\text{avg}}(\lambda) = \mathbf{E}_{\tau} [v(\tau, \lambda)],$$

which is captured in the simulations by clustering real-time prices, and averaging over each cluster.

In these numerical simulations, we average over random instances of price and storage initial states. We set $N = 288$, which corresponds to a period of 24 hours, where real-time

prices are updated once in every 5 minutes. The storage system implements the optimal policy given in Theorem 1. For the purpose of these computations, we use the same price distribution for all k (i.e. we assume that prices are independently and identically distributed). For generating random price sequences, we simulate a discretized truncated log-normal distribution with $\bar{\lambda} = 52$ and $\sigma = 22$, because the log-normal distribution qualitatively resembles the empirical distribution of prices from ISO NE and PJM, at least for the cases that we tested. Also, a mean of 52 and standard deviation of 22 are realistic choices for real-time energy prices in markets such as PJM and ISO NE on a day with moderate volatility. Based on our results in Section 4.2, for these model parameters, a storage capacity of $\bar{s} = 5\bar{v}$ is a reasonable choice for all consumers. We set \hat{t} equal to $\bar{\lambda}$ in these simulations, just like all previous simulations. We also set $\bar{v} = 10$, and $h_k^i = 0$ for all $k \leq N$ and $i < n$.

Figure 6 illustrates how the average storage response changes as a function of price.

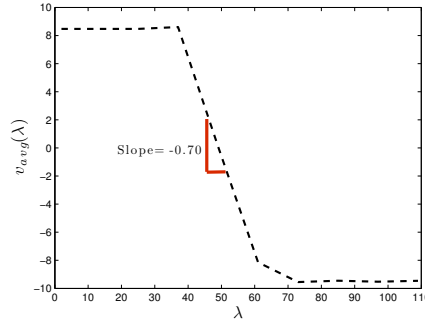


Figure 6 Average storage response vs. price, using the discretized truncated log-normal price distribution.

As the plot suggests, the expected demand seems to be considerably more responsive to changes in the prices that fall in the mid-portion of the price range. This portion serves as a steep transition region, in which the policy quickly switches from the “buy” policy to the “sell” policy.

To characterize price elasticity, let us first recall the standard definition of PED:

$$PED = \frac{\Delta d/d}{\Delta \lambda/\lambda} \quad (22)$$

where d denotes demand. To characterize PED more accurately, one needs to bear in mind that the overall PED should have the firm component of demand in it. Hence, in this setup, we set $d = d^f + v_{avg}(\lambda)$, where d^f denotes the firm component of demand. We can observe in Figure 6 that the average PED is almost zero for prices that are considerably larger or smaller than the mean price, and only in the mid-portion of the plot (i.e. around the mean price) we notice a substantial average PED. One can verify using equation (22) with

Table 3 Average price elasticity of demand around the mean price using different values for d^f

d^f	Average Price Elasticity
\bar{v}	-3.6
$3\bar{v}$	-1.2
$8\bar{v}$	-0.45

$d = d^f + v_{avg}(\lambda)$ that the average PED (i.e. the PED computed using the average storage response) depends on how the average storage response compares with the firm demand. Table 3 shows the average PED around the mean price, using different values for d^f .

5.2. Price Responsiveness of a Storage System

In this subsection, we present a computational framework for understanding the behavior of storage as a function of price and the amount of stored energy, and for characterization of the buy/sell phase transition region in the price-state plane. In order to eliminate stage dependence, we will consider the infinite-horizon version of the storage problem (5) and perform policy iteration (see, e.g., Bertsekas 2000) to numerically obtain a stationary (stage-independent) policy for purchase/sale as a function of both state and price. This section provides a qualitative picture of the structural characteristics of the behavior of storage, and a framework for estimating the PED as a function of the state. Herein, we will assume that the consumer starts with an empty storage, implying that the states would only take on integer multiples of the ramp constraint. We also use the same price distribution for all stages (i.e. we assume that prices are i.i.d). In our computations, we first use a discretized truncated log-normal price distribution, and then compare the results against the case of a discretized uniform distribution. For both distributions, we use a mean of about $\bar{\lambda} = 50$ and a standard deviation of about $\sigma = 30$, and also the same support. We set $\bar{v} = 1$ and $n = 10$. Figure 7 illustrates how the storage response varies with price for three cases of the state (when the storage is empty ($i = 0$), when the storage is half full ($i = n/2$), and when the storage is nearly full ($i = n - 1$)).

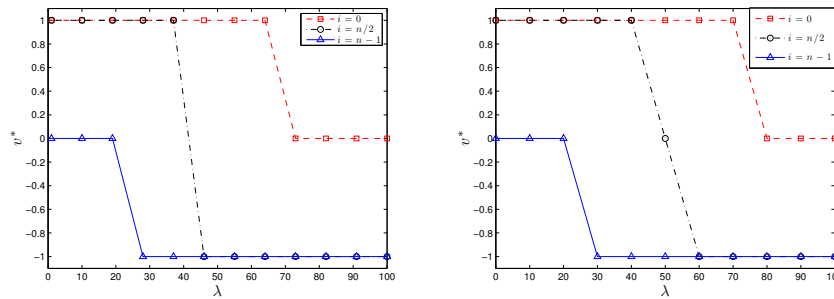


Figure 7 Storage response vs. price for three sample states for discretized truncated log-normal (left) and discretized uniform (right) price distributions, both with mean 50

Note in Figure 7 the considerable effect of the storage state on storage response, compared to the small effect of the price distribution. More specifically, for both distributions, when the storage is empty, the optimal policy recommends purchasing from the grid even when the prices are somewhat above the mean price. Note that for the log-normal case, this policy change occurs at a slightly lower price because of the left skewness of the log-normal distribution. Though, when the storage is half full, we switch from the “buy-it-all” policy to the “sell-it-all policy” right at the mean price for the uniform distribution; for the log-normal distribution, this policy change occurs slightly before the mean price, which is again due to the left skewness of the log-normal distribution. Finally, when the storage is nearly full, for both distributions the optimal policy is to sell as much energy as possible for most prices, and to do nothing for the low prices.

The transition points in the infinite-horizon policy from sell-it-all to buy-it-all on the $s - \lambda$ plane are shown in Figure 8. Any point to the left of and/or below the transition points is a buying policy, which corresponds to $v^*(s, \lambda) = \bar{v}$, and any point to the right of and/or above the transition points is a selling policy, which corresponds to $v^*(s, \lambda) = -\bar{v}$. The plus signs show a direct transition from buying to selling when moving along the vertical axis, i.e., as storage state varies, unless they are immediately followed by a star on their right. The stars denote a transition through a “Do Nothing” policy when moving along the horizontal axis, i.e., as price varies. Therefore, at the prices denoted by * we have $v^* = 0$. Figure 8 clearly illustrates the interplay between the state-dependence and

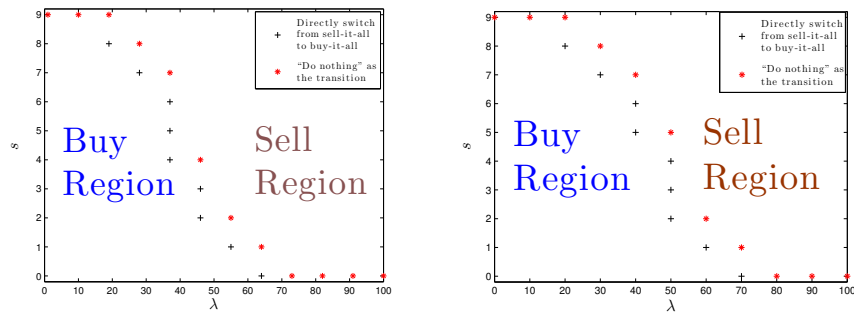


Figure 8 Occurrence points of policy change from sell-it-all to buy-it-all for discretized truncated log-normal (left) and discretized uniform (right) price distributions

the price-dependence of the storage response. Note also in Figure 8 the small effect of the price distribution on the storage response. The optimal policy for the log-normal case is slightly more shifted to left compared to that of the uniform case, which is again caused by the left skewness of the log-normal distribution.

We can now use the above results to characterize the PED as a function of the storage state. In order to compute the overall PED using (22), once again we quantify the demand (d) such that it includes the firm component of the demand as well:

$$d = d^f + v^*(s, \lambda). \quad (23)$$

We have $PED(s) = 0$ for all the points in the “Buy Region” and in the “Sell Region” because the storage response is constant in those regions. However, around the transition curve, the PED is non-zero because $\Delta d < 0$ in that region. But Δd can only take two values: $\Delta d = -2\bar{v}$ when there is a direct transition from the “Buy Region” to the “Sell Region”, and $\Delta d = -\bar{v}$ when the transition is through a “Do Nothing” policy. Hence, we have

$$PED(s) = -\frac{2\bar{v}\lambda}{(d^f + v^*(s, \lambda))\Delta\lambda} \quad \text{and} \quad PED(s) = -\frac{\bar{v}\lambda}{(d^f + v^*(s, \lambda))\Delta\lambda}$$

around the points denoted by $+$ and $*$, respectively, where, depending on the point at which we choose to compute PED, v^* takes on one of the values in $\{-\bar{v}, 0, \bar{v}\}$.

5.3. Impact of Market-Based Operation of Storage on the Required Reserves

In this subsection, we will evaluate how the optimal response from storage affects the amount of reserves (both generation and demand) required to guarantee that supply and demand can be matched. We assume that the Independent System Operator (ISO) primarily uses renewable generation with zero marginal cost and a conventional generation source with a quadratic cost to meet the demand. It is assumed that the overall demand consists of the storage response and a deterministic demand with both elastic and inelastic components. In our simulated setup, at the beginning of each pricing period, the ISO predicts the amount of renewable generation available during that period, possibly with some error. We assume that the amount of renewable generation in each period is i.i.d and is sampled from the bimodal distribution shown in Figure 9, which can for instance, correspond to a system with two renewable sources, one with high capacity and one with low capacity. In the event that during any time period, there is a shortfall in renewable generation compared to what was predicted, the ISO extracts from the generation reserves. Similarly, if there is an excess of renewable generation, the ISO would direct the excess generation to the demand reserves. We assume that the consumers can learn only the stationary distribution of prices (under the assumption of i.i.d prices), and not the exact mechanism by which prices are generated, and that the consumers are rational, and hence, they have no incentive to manipulate their storage response. Thus, the storage management policy discussed in Section 5.2 is deemed optimal by the consumers. The setup is

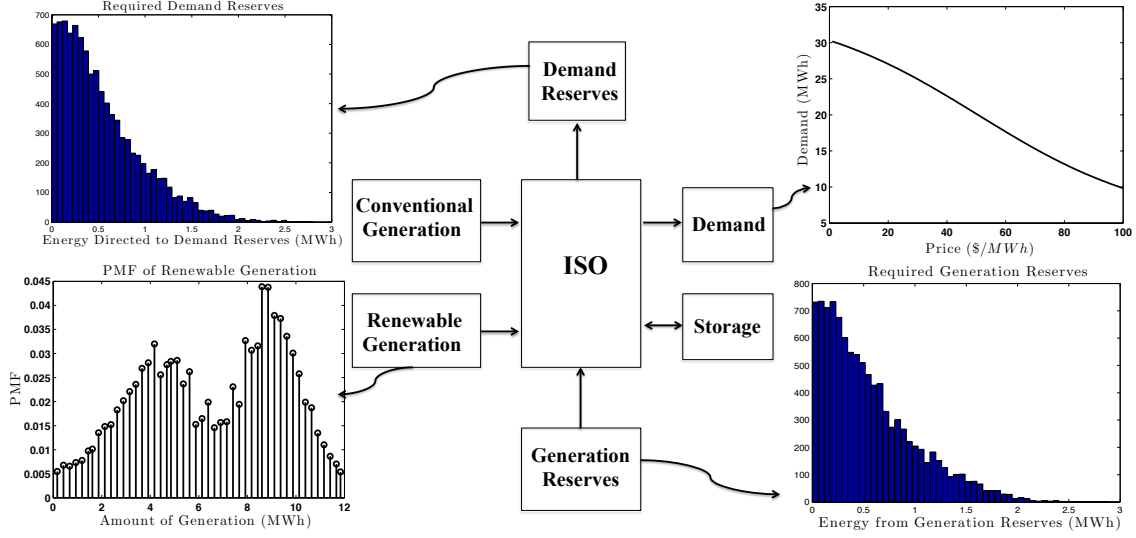


Figure 9 Matching demand and supply by the ISO

shown in Figure 9. In this setup, at the beginning of each time period, the ISO determines the clearing price by solving the following equation for λ_k :

$$a\lambda_k + \hat{w}_k = v^*(s, \lambda_k) + d_k(\lambda_k) \quad (24)$$

where \hat{w}_k is the predicted amount of renewable generation during the k -th time period, $v^*(s, \lambda_k)$ is the optimal response of storage obtained using policy iteration as discussed in Section 5.2, $d_k(\lambda_k)$ is the deterministic demand, and $a\lambda_k$ is the optimal response of a supplier with quadratic cost $c(x) = \frac{1}{2a}x^2$, to a given price λ_k . For the purpose of numerical computations, we assume that $d_k(\lambda_k)$ is a logistic function (as shown in Figure 9)

$$d_k(\lambda_k) = 5 \text{ MWh} + \frac{(30 \text{ MWh})e^{-\lambda_k/30}}{e^{-\lambda_k/30} + e^{-\bar{\lambda}/30}}. \quad (25)$$

The choice of a logistic function for characterizing the deterministic demand was due to the fact that qualitatively similar functions have been used in the past for modeling price responsive demand in electricity markets (see, e.g., Carrión, Conejo, and Arroyo 2007), this choice is otherwise arbitrary and not central to our analysis or conclusions. The choice of the numerical values of the parameters of the logistic function is reasonable for the hourly load in a small power system. For all k we set $\bar{s}/\bar{v} = n = 5$, the average price $\bar{\lambda} = 50$ \$/MWh, and vary the ramp constraint for experimentation, looking at three cases: $\bar{v} = 0.25$ MW, $\bar{v} = 0.5$ MW, and $\bar{v} = 1$ MW such that the storage capacity \bar{s} is roughly 4.17%, 8.33%, and 16.67%, respectively, of the maximum possible demand in the absence of storage. We also set $a = 4$ (MWh)²/\$ so that the average price in the simulated setup is about 50 \$/MWh, which is close to the typical average hourly prices in various electricity markets. Note that since we are fixing n , all the storage devices across the grid become synchronized in the

Table 4 Change in the required reserves for various values of ramp and reliability levels

$\bar{s}/\text{max hourly dem}$	Reliability	Change in demand reserves	Change in generation reserves
$\frac{1}{24}$	100%	7.0% \uparrow	9.9% \uparrow
	99%	3.1% \uparrow	2.7% \uparrow
	98%	2.4% \uparrow	3.8% \uparrow
$\frac{1}{12}$	100%	22% \uparrow	24% \uparrow
	99%	13% \uparrow	15% \uparrow
	98%	12% \uparrow	14% \uparrow
$\frac{1}{6}$	100%	40% \uparrow	42% \uparrow
	99%	43% \uparrow	45% \uparrow
	98%	45% \uparrow	46% \uparrow

steady state, and hence, we can use one large value for \bar{v} which serves as a proxy for the cumulative ramp rate constraint across the grid.

We further assume that the prediction of the amount of renewable generation available in each period has an error that is sampled from a truncated normal distribution with zero mean and standard deviation 0.5 MWh, and that the ISO estimates the storage state with an error that has a discrete uniform distribution over the support $\{-\bar{v}, 0, \bar{v}\}$. We compute the percentage change in the amount of reserves needed in the presence of storage compared to the case of no storage, for each storage capacity mentioned above. In each scenario, we first examine the 100% reliability level (i.e. when demand and supply are guaranteed to match roughly 100% of the time) and then we repeat the computations for 99% and 98% reliability levels.

As the simulation results in Table 4 suggest, market based operation of storage may require an increase in the amount of required reserves (both demand and supply reserves). The larger the amount of energy that can be stored, or extracted from storage, the higher the amount of demand and supply reserves needed. Note that in the largest case of storage reported in Table 4, the storage capacity is only one sixth of the maximum demand, while the reserves need to be expanded by about 40% to accommodate the integration of this much storage (the histograms of the required reserves for this case of storage capacity are shown in Figure 9).

This effect is mainly caused by the threshold-based, state dependent response of storage as discussed in Section 5.2 (Figure 8). When the prices are high due to lack of renewable generation, the storage system does not necessarily respond to the price incentive the way the ISO would want it to, because even if the prices are above the average, the optimal policy for the storage could be to buy if the storage level is low. Hence, in the event that the actual amount of renewable generation falls short of what was predicted, if the ISO

also makes an error in estimating the storage state, a considerable amount of energy may be extracted from the grid by the storage device, which would force the ISO to use the generation reserves even more. Analogously, if there is an excess of renewable generation, the ISO may be forced to use the demand reserves even more if the storage device feeds energy back into the grid. A higher penetration of storage, and hence, a higher aggregate ramp rate amplifies this issue because the impact of erroneous state estimations would be even more profound. However, the generation reserves are typically fast generators with high economic and environmental costs. Hence, the ISO may need to regulate market-based operation of storage devices to prevent such impacts, and/or design mechanisms for pricing these externalities. Also, the ISO may need to design mechanisms that guarantee its access to the exact value of the storage state, because the intense interplay between state and price is such that estimating the storage state even with a small error can heavily impact the storage response.

6. Conclusion

In this paper, we proposed a dynamic model for optimal control of storage under ramp constraints and exogenous, stochastic prices. We derived the associated optimal policy and value function, and gave explicit formulas for their computation. Moreover, we derived an analytical upperbound on the long-term average economic value of storage, which is valid for any achievable realization of prices over a fixed support, and highlights the dependence of the value on ramp constraints and capacity. This result can be useful in assessing viability of investment in electricity storage. We also showed that while the value of storage is a non-decreasing function of price volatility, due to finite ramping rates, the value of storage saturates quickly as capacity increases, regardless of price volatility. We highlighted the dependence of the response of storage to prices on the internal state of storage, and also, by averaging out state and stage dependence, we showed that in expectation, storage induces a considerable amount of price elasticity near the average price. We also showed that if the ISO does not have perfect information about the exact value of the storage state, the reserves may need to be expanded to accommodate market-based operation of storage.

Our results provide insight into learning the behavior of storage, particularly modeling and estimating the response of a ramp-constrained storage system when used as an arbitrage mechanism. We used price data from real-time wholesale markets to examine the sensitivity of the optimal policy and the value of storage to our assumption of independent prices. The relatively high competitive ratios that we found when testing the optimal policy suggest that the sensitivity of the policy to the assumption of independent prices is

low. The competitive ratios found in this paper by applying the optimal policy can be used as a benchmark for comparison with other results that assume more complicated models of the stochastic price process.

We did not directly consider the effects of inefficiencies such as conversion losses and aging on the optimal policy. While such inefficiencies certainly limit the economic value of storage, an important question to ask is what is their effect on the response of storage to prices? For the case of battery storage, aging is proportional to the amount of current withdrawn or injected. Intuitively, this would make buying energy for, or selling energy from storage less profitable at moderate prices. Within the class of threshold policies, the corresponding optimal selling thresholds would be higher and the buying thresholds would be lower than what we have derived in this paper and the “do nothing” range would be wider. Qualitatively, this would mean an overall choppier response from storage, with high elasticity over a narrow range of prices and low elasticity over a wide range, which would be undesirable from a system operation and reliability point of view.

While our paper largely focused on the economic value of storage, it is important to recognize and quantify the environmental and the reliability value of storage. With proper control policies, storage can help matching stochastic supply with demand, improving system frequency and voltage profiles, and possibly mitigating large blackouts. The development of a systematic framework for quantifying the value of storage, the trade-offs between reliability, environmental, and economic value of storage, and the design of real-time pricing and market mechanisms for optimally striking these trade-offs are fundamental and important directions for future research.

References

- A123 Systems. (2012). Smart Grid Stabilization and Smart Energy Storage Solutions. www.a123systems.com. Retrieved in April, 2012.
- Bannister, C. H., & Kaye, R. J. (1991). A rapid method for optimisation of linear systems with storage. *Operations Research*, 39, 220-232.
- Bellman, R. (1956). On the theory of dynamic programming a warehousing problem. *Management Science*, 2, 272-275.
- Bertsekas, D. P. (2000). *Dynamic Programming and Opt. Control - Vols. I & II*. (2nd ed.). Belmont: Athena Scientific.
- Bitar, E., Giani, A., Rajagopal, R., Varagnolo, D., Khargonekar, P., Poolla, K., & Varaiya, P. (2010). Optimal contracts for wind power producers in electricity markets. *Proceedings of 49th IEEE Conference on Decision and Control*, Atlanta, GA.

-
- Brown, P. D., Lopes J. A. P., & Matos, M. A. (2008). Optimization of pumped storage capacity in an isolated power system with large renewable penetration. *IEEE Trans. on Power Systems*, 23, 523-53.
- Cahn, A. S. (1948). The warehouse problem. *Bulletin of the American Mathematical Society*, 54, 1073.
- Carrión, M., Conejo, A. J., & Arroyo, J. M. (2007). Forward contracting and selling price determination for a retailer. *IEEE Transactions on Power Systems*, 22, 2105-2114.
- Charnes, A., Drèze, J., & Miller, M. (1966). Decision and horizon rules for stochastic planning problems: A linear example. *Econometrica*, 34, 307-330.
- Devalkar, S., Anupindi, R. , & Sinha, A. (2011). Integrated optimization of procurement, processing and trade of commodities. *Operations Research*, 59, 1369-1381.
- Dreyfus, S. E. (1957). An analytic solution of the warehouse problem. *Management Science*, 4, 99-104.
- Drouin, N., Gautier, A., Lamond, B. F., & Lang, P. (1996). Piecewise affine approximations for the control of a one-reservoir hydroelectric system. *European Journal of Operational Research*, 89, 53-69.
- Faghih, A., Roozbehani, M., & Dahleh, M. A. (2011). Optimal utilization of storage and the induced price elasticity of demand in the presence of ramp constraints. *Proceedings of the 50th IEEE Conference on Decision and Control*, Orlando, FL.
- Faruqui, A., & George, S. S. (2002). The value of dynamic pricing in mass markets. *The Electr. Jour*, 15, 45-55.
- Federgruen, A., & Zipkin, P. H. (1986). An inventory model with limited production capacity and uncertain demands II: The discounted-cost criterion. *Mathematics of Operations Research*, 11, 208-215.
- Goel, A., & Gutierrez, G. (2009). Integrating commodity markets in the optimal procurement policies of a stochastic inventory system. Working Paper, University of Texas, Austin, TX.
- Gonzalez, J. G., de la Muela, R. M. R., Santos, L. M., & Gonzalez, A. M. (2008). Stochastic joint optimization of wind generation and pumped-storage units in an electricity market. *IEEE Transactions on Power Systems*, 23, 460-468.
- Harsha, P., & Dahleh, M. A. (2011). Optimal sizing of energy storage for efficient integration of renewable energy. *Proceedings of the 50th IEEE Conference on Decision and Control*, Orlando, FL.
- ISO New England website. (2011). Historical hourly price data. <http://www.iso-ne.com>. Retrieved in Dec, 2011.
- Kaminski, V., Feng, Y., & Pang, Z. (2008). Value, trading strategies and financial investment of natural gas storage assets. *Northern Finance Association Conference*, Kananaskis, Canada.
- Kapuscinski, R., & Tayur, S. (1998). Capacitated production-inventory model with periodic demand. *Operations Research*, 46, 899-911.
- Kim, J. H., & Powell, W. B. (2011). Optimal energy commitments with storage and intermittent supply. *Operations Research*, 59, 1347-1360.

- Kirschen, D. S., Strbac, G., Cumperayot, P., & de Paiva Mendes, D. (2000). Factoring the elasticity of demand in electricity prices. *IEEE Trans on Power Systems*, 15, 612-617.
- Kirschen, D. S. (2003). Demand-side view of electricity markets. *IEEE Trans on Power Systems*, 18, 520-527.
- Korpaas, M., Holen, A. T., & Hildrum, R. (2003). Operation and sizing of energy storage for wind power plants in a market system. *International Journal of Electrical Power & Energy Systems*, 25, 599-606.
- Lamond, B. F., Monroe, S. L., & Sobel, M. J. (1995). A reservoir hydroelectric system: Exactly and approximately optimal policies. *European Journal of Operational Research*, 81, 535-542.
- Lee, T. Y., & Chen, N. (1995). Determination of optimal contract capacities and optimal sizes of battery energy storage systems for time-of-use rates industrial customers. *IEEE Trans on Energy Conversion*, 10, 562-568.
- PJM Interconnection website. (2011). Data Dictionary: Monthly Real-Time Locational Marginal Pricing. <http://www.pjm.com>. Retrieved in Dec, 2011.
- Qian, H., Zhang, J., Lai, J. S., & Yu, W. (2011). A high-efficiency grid-tie battery energy storage system. *IEEE Trans. on Power Electronics*, 26, 886-896.
- Rempala, R. (1994). Optimal strategy in a trading problem with stochastic prices. *System Modelling and Optimization, Lecture Notes in Control and Information Sciences*, 197, 560-566.
- Roozbehani, M., Dahleh, M. A., & Mitter, S. (2011). Volatility of power grids under real-time pricing. *IEEE Transactions on Power Systems*, Accepted subject to revisions. Available at <http://mit.edu/mardavij/www/publications.html>
- Secomandi, N. (2010). Optimal commodity trading with a capacitated storage asset. *Management Science*, 56, 449-467.
- Wu, O., Wang, D., & Qin, Z. (2011). Seasonal energy storage operations with limited flexibility. Working paper, Ross School of Business, University of Michigan, Ann Arbor, MI.
- Yusta, J. M., & Dominguez, J. A. (2002). Measuring and modeling of industrial demand response to alternative prices of the electricity. *Proceedings of Power Systems Computation Conference*, Sevilla, Spain.

7. Appendix

Proof of Theorem 1:

Proof. This proof proceeds by induction. Let us for the moment assume that the value function $V_k(\cdot) \stackrel{\text{def}}{=} \mathbf{E}[J_k(\cdot)]$ has the form defined in (12).

From the dynamic programming algorithm, for $k < N$, we have

$$J_k(s_k) = h_k(s_k) + \min_{v_k \in [\max(-s_k, -\bar{v}), \bar{v}]} \{ \lambda_k v_k + \mathbf{E}[J_{k+1}(s_k + v_k)] \}, \quad (\text{A.1})$$

where the penalty functions $h_k(s_k)$ are as defined in (9).

Then, we use the general form in (12) for $\mathbf{E}[J_k(s_k)]$ and the state evolution equation in (2), and let $i \in \bar{\mathbb{Z}}_+$ be such that $v_k + s_k \in [i\bar{v}, (i+1)\bar{v})$. Applying the induction step to (A.1) we obtain:

$$J_k(s_k) = h_k(s_k) + \min_{v_k \in [\max(-s_k, -\bar{v}), \bar{v}]} \{ \lambda_k v_k - t_{k+1}^i(s_k + v_k) + e_{k+1}^i \} \quad (\text{A.2})$$

Solving the optimization problem in (A.2) yields the optimal policy shown in (10). Now we can plug in the optimal policy to obtain $J_k(s_k)$. Let $i \in \bar{\mathbb{Z}}_+$ be such that $s_k \in [i\bar{v}, (i+1)\bar{v})$. If $0 \leq s_k < \bar{v}$, then

$$J_k(s_k) = \begin{cases} (h_k^0 - \lambda_k)s_k + e_{k+1}^0 + c_k^0 & \text{if } t_{k+1}^0 < \lambda_k \\ (\lambda_k - t_{k+1}^1)\bar{v} + (h_k^0 - \lambda_k)s_k + e_{k+1}^1 + c_k^0 & \text{if } t_{k+1}^1 < \lambda_k \leq t_{k+1}^0 \\ (\lambda_k - t_{k+1}^1)\bar{v} + (h_k^0 - t_{k+1}^1)s_k + e_{k+1}^1 + c_k^0 & \text{if } \lambda_k \leq t_{k+1}^1 \end{cases}$$

and if $s_k \geq \bar{v}$, then

$$J_k(s_k) = \begin{cases} -(\lambda_k - t_{k+1}^{i-1})\bar{v} + (h_k^i - t_{k+1}^{i-1})s_k + e_{k+1}^{i-1} + c_k^i & \text{if } t_{k+1}^{i-1} < \lambda_k \\ (\lambda_k - t_{k+1}^i)\bar{v} + (h_k^i - \lambda_k)s_k + e_{k+1}^i + c_k^i & \text{if } t_{k+1}^i < \lambda_k \leq t_{k+1}^{i-1} \\ (\lambda_k - t_{k+1}^{i+1})(i+1)\bar{v} + (h_k^i - \lambda_k)s_k + e_{k+1}^{i+1} + c_k^i & \text{if } t_{k+1}^{i+1} < \lambda_k \leq t_{k+1}^i \\ (\lambda_k - t_{k+1}^{i+1})\bar{v} + (h_k^i - t_{k+1}^{i+1})s_k + e_{k+1}^{i+1} + c_k^i & \text{if } \lambda_k \leq t_{k+1}^{i+1} \end{cases}$$

Let us recall that for a piecewise function f defined according to

$$f(x) = \begin{cases} f_1(r)x + g_1(r) & \text{if } r \leq \alpha \\ f_2(r)x + g_2(r) & \text{if } r > \alpha \end{cases}$$

we have

$$\begin{aligned} \mathbf{E}[f(x)] &= \mathbf{E}[f(x) | r \leq \alpha] \mathbf{P}(r \leq \alpha) + \mathbf{E}[f(x) | r > \alpha] \mathbf{P}(r > \alpha) \\ &\quad + \mathbf{E}[g_1(r) | r \leq \alpha] \mathbf{P}(r \leq \alpha) + \mathbf{E}[g_2(r) | r > \alpha] \mathbf{P}(r > \alpha) \end{aligned} \quad (\text{A.3})$$

Also, note that:

$$\mathbf{E}[f(r) | r \leq \alpha] \mathbf{P}(r \leq \alpha) = \sum_{\theta=r_{\min}}^{\alpha} f(r) P_R(\theta) \quad (\text{A.4})$$

Now, let us apply (A.3) and (A.4) to the equations derived above for $J_k(s_k)$, and compute $\mathbf{E}[J_k(s_k)]$ for $k < N$, which leads to the following results:

if $0 \leq s_k < \bar{v}$, then

$$\mathbf{E}[J_k(s_k)] = e_k^0 - s_k[t_{k+1}^1 F_k(t_{k+1}^1) - h_k^0] + \sum_{t_{k+1}^1 < \theta \leq \lambda_k^{\max}} \theta P_k(\theta)$$

if $s_k \geq \bar{v}$, then

$$\mathbf{E}[J_k(s_k)] = e_k^i - s_k[t_{k+1}^{i-1}(1 - F_k(t_{k+1}^{i-1})) + t_{k+1}^{i+1} F_k(t_{k+1}^{i+1}) - h_k^i] + \sum_{t_{k+1}^{i+1} < \theta \leq t_{k+1}^{i-1}} \theta P_k(\theta)$$

where e_k^i denotes the sum of the terms that have not been multiplied by s_k . Hence, the thresholds for $k < N$ and $i \in \mathbb{Z}_+$ are given by:

$$\begin{aligned} t_k^0 &= t_{k+1}^1 F_k(t_{k+1}^1) - h_k^0 + \sum_{t_{k+1}^1 < \theta \leq \lambda_k^{\max}} \theta P_k(\theta) \\ t_k^i &= t_{k+1}^{i-1}(1 - F_k(t_{k+1}^{i-1})) + t_{k+1}^{i+1} F_k(t_{k+1}^{i+1}) - h_k^i + \sum_{t_{k+1}^{i+1} < \theta \leq t_{k+1}^{i-1}} \theta P_k(\theta) \end{aligned}$$

After some algebra, the above results can be written in the form presented in (11). The derivation of the intercepts e_k^i follows the same procedure.

The next step is to verify, using induction, that the thresholds at each stage (i.e. t_k^i) are non-increasing functions of i . Considering that t_{k-1}^0 has a different general form than t_{k-1}^i for $i > 0$, we first need to show that $t_{k-1}^1 \leq t_{k-1}^0$ assuming that $t_k^{i+1} \leq t_k^i$ for all i . Hence, we need to first show that

$$t_k^1 F_{k-1}(t_k^1) - h_{k-1}^0 + \sum_{t_k^1 < \theta \leq \lambda_{k-1}^{\max}} \theta P_{k-1}(\theta) \geq t_k^0 (1 - F_{k-1}(t_k^0)) + t_k^2 F_{k-1}(t_k^2) - h_{k-1}^1 + \sum_{t_k^2 < \theta \leq t_k^0} \theta P_{k-1}(\theta)$$

Knowing that $h_{k-1}^1 \geq h_{k-1}^0$, we can remove $-h_{k-1}^1$ and $-h_{k-1}^0$ from both sides. Then, by writing the cumulative distribution functions as summations, and rearranging some terms, the above can be rewritten as:

$$t_k^0 \sum_{\lambda_{k-1}^{\min} \leq \theta \leq t_k^0} P_{k-1}(\theta) + \sum_{\lambda_{k-1}^{\min} \leq \theta \leq t_k^1} P_{k-1}(\theta) t_k^1 + \sum_{t_k^1 < \theta \leq \lambda_{k-1}^{\max}} \theta P_{k-1}(\theta) \geq t_k^0 + \sum_{\lambda_{k-1}^{\min} \leq \theta \leq t_k^2} P_{k-1}(\theta) t_k^2 + \sum_{t_k^2 < \theta \leq t_k^0} \theta P_{k-1}(\theta)$$

By breaking each summation into disjoint intervals, and factoring all the terms in the same interval and merging them into one summation, the above can be rewritten as:

$$t_k^0 \leq \sum_{\lambda_{k-1}^{\min} \leq \theta \leq t_k^2} P_{k-1}(\theta) (t_k^0 + t_k^1 - t_k^2) + \sum_{t_k^2 < \theta \leq t_k^1} P_{k-1}(\theta) (t_k^0 + t_k^1 - \theta) + \sum_{t_k^1 < \theta \leq t_k^0} P_{k-1}(\theta) t_k^0 + \sum_{t_k^0 < \theta \leq \lambda_{k-1}^{\max}} \theta P_{k-1}(\theta)$$

We can see that in the equation on the right hand side (RHS) of the inequality shown

above, all the terms that have been multiplied by $P_{k-1}(\theta)$ inside the summations are greater than or equal to t_k^0 ; we also know that $\sum_{\theta=\lambda_{k-1}^{\min}}^{\lambda_{k-1}^{\max}} P_{k-1}(\theta) = 1$. Hence, we verify by inspection that the RHS equation in the inequality shown above is always greater than or equal to t_k^0 .

Finally, we have assigned a non-positive cost of $-\hat{t}$ to each unit of energy left in storage at stage N (i.e. $\mathbf{E}[J_N(s_N)] = -\hat{t}s_N$ for $s_N \leq \bar{s}$). Also, a very small (negative) value is assigned to the thresholds for $i \geq n$ at stage N , to make sure we will not exceed the storage capacity in this stage. Hence, t_N^i is a non-increasing function of i , and $t_N^{i+1} \leq t_N^i$ is satisfied.

We must now verify that $t_{k-1}^{i+1} \leq t_{k-1}^i$ for $i \in \mathbb{Z}_+$ assuming that t_k^i is a non-increasing function of i . So, we need to verify:

$$\begin{aligned} t_k^i (1 - F_{k-1}(t_k^i)) + t_k^{i+2} F_{k-1}(t_k^{i+2}) - h_{k-1}^{i+1} + \sum_{t_k^{i+2} < \theta \leq t_k^i} \theta P_{k-1}(\theta) &\leq t_k^{i-1} (1 - F_{k-1}(t_k^{i-1})) \\ &\quad + t_k^{i+1} F_{k-1}(t_k^{i+1}) + \sum_{t_k^{i+1} < \theta \leq t_k^{i-1}} \theta P_{k-1}(\theta) - h_{k-1}^i \end{aligned}$$

Knowing that $h_{k-1}^{i+1} \geq h_{k-1}^i$, we can remove $-h_{k-1}^{i+1}$ and $-h_{k-1}^i$ from both sides. Then, by writing the cumulative distribution functions as summations, and rearranging some terms, the above can be rewritten as:

$$\begin{aligned} t_k^i + t_k^{i-1} \sum_{\lambda_{k-1}^{\min} \leq \theta \leq t_k^{i-1}} P_{k-1}(\theta) + t_k^{i+2} \sum_{\lambda_{k-1}^{\min} \leq \theta \leq t_k^{i+2}} P_{k-1}(\theta) + \sum_{t_k^{i+2} < \theta \leq t_k^i} \theta P_{k-1}(\theta) &\leq t_k^{i-1} + t_k^i \sum_{\lambda_{k-1}^{\min} \leq \theta \leq t_k^i} P_{k-1}(\theta) \\ &\quad + t_k^{i+1} \sum_{\lambda_{k-1}^{\min} \leq \theta \leq t_k^{i+1}} P_{k-1}(\theta) + \sum_{t_k^{i+1} < \theta \leq t_k^{i-1}} \theta P_{k-1}(\theta) \end{aligned}$$

By taking all the summations to the right hand side and taking t_k^i to the left hand side of the inequality, breaking each summation into disjoint intervals, and factoring all the terms in the same interval and merging them into one summation, the above can be rewritten as:

$$\begin{aligned} t_k^{i-1} - t_k^i &\geq \sum_{\lambda_{k-1}^{\min} \leq \theta \leq t_k^{i+2}} P_{k-1}(\theta) (t_k^{i-1} - t_k^i - (t_k^{i+1} - t_k^{i+2})) + \sum_{t_k^{i+2} < \theta \leq t_k^{i+1}} P_{k-1}(\theta) (t_k^{i-1} - t_k^i - (t_k^{i+1} - \theta)) \\ &\quad + \sum_{t_k^{i+1} < \theta \leq t_k^i} P_{k-1}(\theta) (t_k^{i-1} - t_k^i) + \sum_{t_k^i < \theta \leq t_k^{i-1}} P_{k-1}(\theta) (t_k^{i-1} - \theta) \end{aligned}$$

We can see that in the equation on the RHS of the inequality shown above, all the terms that have been multiplied by $P_{k-1}(\theta)$ inside the summations are less than or equal to $t_k^{i-1} - t_k^i$; we also know that $\sum_{\theta=\lambda_{k-1}^{\min}}^{t_k^{i-1}} P_{k-1}(\theta) \leq 1$. So, given that the summations in the RHS do not overlap, we can verify by inspection that the RHS equation in the inequality shown above is always less than or equal to $t_k^{i-1} - t_k^i$.

The last step is to show that the value function is a continuous function. This can be done by induction. We have defined $\mathbf{E}[J_N(s_N)]$ in such a way that it is convex. We also have defined $h_k(s_k)$ to be convex for all k . Looking at the equations of $J_k(s_k)$ in (A.1) and (A.2), given that $\mathbf{E}[J_{k+1}(s_{k+1})]$ is convex, one can observe that the continuity of $J_k(s_k)$ is satisfied because v_k^* is a continuous function of s_k . Similarly, considering the equations obtained for v_k^* , the continuity of $\mathbf{E}[J_k(s_k)]$ for all k follows from the fact that the expectation (i.e. $\mathbf{E}[J_k(s_k)]$) is simply a convex combination of continuous functions. This completes the proof. \square

Proof of Theorem 2:

Proof. The results of Theorem 1 establish that for any finite N , the optimal policy is a threshold policy regardless of the price distribution. It can be verified that the form of the optimal policy extends to the infinite-horizon case in the sense that the infinite-horizon optimal policy is similar to the policy defined in Theorem 1, with stationary thresholds that are only a function of the storage state. Furthermore, the optimal average cost is independent of the initial state. Hence, without loss of generality we may assume that we start from $s_0 = 0$, and therefore, the storage state would only take values that are integer multiples of the ramp rate, i.e., $s \in \{i\bar{v} \mid i \in 0, 1, \dots, n\}$. This assumption simplifies the optimal policy to either buy as much as \bar{v} , or do nothing, or sell as much as \bar{v} . For conciseness and ease of notation, and without loss of generality, we present the remainder of the proof for the case of $\bar{v} = 1$, $\lambda_{\min} = 0$, and $\lambda_{\max} = 1$. The proof for the general case is similar. Let $\bar{t}(s)$ and $\underline{t}(s)$ denote the threshold associated with selling and buying, respectively, at state s . Thus, if the stage price λ_k falls within the interval $(\underline{t}(s), \bar{t}(s))$, the decision is to do nothing and no cost is incurred. Let J_k denote the cost at stage k . Then:

$$\begin{aligned}
\mathbf{E}[J_k] &= \mathbf{E}[\mathbf{E}[J_k|s_k]] = \mathbf{E}[\mathbf{E}[\lambda_k \min\{\bar{s} - s_k, 1\}|s_k, \lambda_k \leq \underline{t}(s_k)]\mathbf{P}(\lambda_k \leq \underline{t}(s_k)) \\
&\quad + \mathbf{E}[\lambda_k \max\{-s_k, -1\}|s_k, \lambda_k \geq \bar{t}(s_k)]\mathbf{P}(\lambda_k \geq \bar{t}(s_k))] \\
&\geq \mathbf{E}[\mathbf{E}[\lambda_k \max\{-s_k, -1\}|s_k, \lambda_k \geq \bar{t}(s_k)]\mathbf{P}(\lambda_k \geq \bar{t}(s_k)))] \\
&\geq -\mathbf{P}(s_k \geq 1)\mathbf{E}[\mathbf{P}(\lambda_k \geq \bar{t}(s_k))] \\
&= -(1 - P(s_k = 0))\mathbf{P}(\lambda_k = 1)
\end{aligned} \tag{A.5}$$

where the first inequality follows from $\lambda_{\min} = 0$, and the second inequality follows from $\lambda_{\max} = 1$ and the fact that selling can happen only when $i \neq 0$. Thus, $\mathbf{E}[J_k]$ is bounded from below by (A.5), with exact equality holding for a two-point distribution with PMF:

$$P_\Lambda(\lambda) = \begin{cases} a & \text{if } \lambda = \lambda_{\max} = 1 \\ 1 - a & \text{if } \lambda = \lambda_{\min} = 0 \\ 0 & \text{otherwise} \end{cases}$$

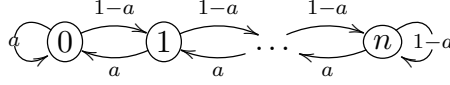


Figure 10 Underlying Markov chain associated with storage state under two-point price distribution

Next, we show that $a = 0.5$ minimizes the stage cost. Under the two-point distribution defined above, the storage state evolves according to the discrete time finite-state Markov chain shown in Figure 10.

Letting p_i denote the steady state probability of state i , it can be verified that $p_0 = 1/(1 + b + \dots + b^n)$, where $b = (1 - a)/a$. Considering that the optimal decision is to either purchase at a price of $\lambda_{\min} = 0$ (which incurs a cost of zero) or sell at a price of $\lambda_{\max} = 1$ (which occurs with probability $a(1 - p_0)$ and incurs a cost of $-\lambda_{\max}\bar{v} = -1$), we have

$$\begin{aligned} \gamma &= \lim_{N \rightarrow \infty} \frac{1}{N} \mathbf{E} \left[\sum_{k=0}^{N-1} \lambda_k v_k \mid s_0 = 0 \right] \\ &= \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \mathbf{E}[\mathbf{E}[\lambda_k \max\{-s_k, -1\} \mid \lambda_k \geq \bar{t}(s_k), s_k] \mathbf{P}(\lambda_k \geq \bar{t}(s_k))] \\ &= - \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \mathbf{P}(s_k \geq 1) \mathbf{P}(\lambda_k = 1) = - \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} (1 - p_0)a, \end{aligned}$$

which yields

$$\gamma = - \frac{b(1 + b + \dots + b^{n-1})}{(b + 1)(1 + b + \dots + b^n)} \quad (\text{A.6})$$

Solving for b to minimize γ yields $b = 1$, which corresponds to $a = 1/2$. This proves (17), (18), and (21). It remains to show that the differential cost function corresponding to (17) is of the form (19).

Let $H(s)$ be the value of the right hand side (RHS) of equation (16) obtained by plugging the proposed solution for $H^*(s)$ and γ^* into equation (16); then we have

$$\begin{aligned} H(s) &= \mathbf{E} \left[\min_{v \in [\max(-s, -\bar{v}), \min(\bar{v}, \bar{s} - s)]} \lambda v - \frac{(i + 1)\lambda_{\min} + (n - i)\lambda_{\max}}{n + 1} (s + v) - \frac{i(i + 1)(\lambda_{\max} - \lambda_{\min})\bar{v}}{2(n + 1)} \right] \\ &\quad + \frac{\bar{v}(\lambda_{\max} - \lambda_{\min})}{2} \frac{n}{n + 1}, \end{aligned} \quad (\text{A.7})$$

where $i\bar{v} \leq s + v < (i + 1)\bar{v}$, for all $i \in \{0, 1, 2, \dots, n - 1\}$.

Therefore, we need to show that $H(s)$, obtained from solving the optimization problem in (A.7), is indeed the $H^*(s)$ defined in (19). It can be verified that the solution to the optimization problem in (A.7) is as follows:

$$v^* = \begin{cases} \max(-s, -\bar{v}) & \text{if } \lambda = \lambda_{\max} \\ \min(\bar{v}, \bar{s} - s) & \text{if } \lambda = \lambda_{\min} \end{cases} \quad (\text{A.8})$$

It can be verified by inspection that if we plug (A.8) into (A.7), the $H(s)$ obtained is indeed the $H^*(s)$ defined in (19). The proof is complete. \square

Proof of Corollary 1:

Proof. Since a two-point distribution with non-zero probability masses placed at the endpoints of the fixed support is always achievable for any $\mu \in (\lambda_{\min}, \lambda_{\max})$, the additional assumption of a fixed mean does not affect the proof of Theorem 2 up to and including equation (A.6). For the proof of this corollary, we need to minimize γ in equation (A.6) with respect to b subject to the constraint $\mu = b\lambda_{\min}/(b+1) + \lambda_{\max}/(b+1)$. This optimization problem has the unique solution $b = (\lambda_{\max} - \mu)/(\mu - \lambda_{\min})$. Recall that $a = 1/(b+1)$. Hence, $a = (\mu - \lambda_{\min})/(\lambda_{\max} - \lambda_{\min})$. This completes the proof. \square